

VLA-Adapter vs. Full Fine-Tuning: Computational Efficiency at Scale on LAVIS

Assignee Research

June 9, 2026

Abstract

This report synthesises findings from 8 peer-reviewed papers addressing the following research question: What is the computational efficiency trade-off between VLA-Adapter and full fine-tuning when scaling to larger multimodal models on the LAVIS benchmark, measured by training time and memory usage. 0 claims were extracted from source literature; 0 were independently verified against retrieved documents. An automated multi-reviewer quality assessment produced a score of 6.5/10. This report is a machine-generated literature synthesis and does not constitute original research.

1 Introduction

This paper examines: Differentially Private Fine-tuning of Language Models. Research question: What is the computational efficiency trade-off between VLA-Adapter and full fine-tuning when scaling to larger multimodal models on the LAVIS benchmark, measured by training time and memory usage?.

2 Methodology

Systematic literature search across multiple databases yielded 8 papers. Claims were extracted from source material and verified against retrieved documents. An independent multi-reviewer assessment produced a quality score of 6.5/10.

3 Results

8 papers retrieved. 0 claims extracted; 0 independently verified. Quality review score: 6.5/10.

4 Limitations

This report is a machine-generated literature synthesis and does not constitute original research. Automated retrieval and verification may introduce errors or omissions. Review scores reflect automated assessment, not human peer review. Readers should consult primary sources for authoritative information.

References

- <http://arxiv.org/abs/2602.08239v1>
- <http://arxiv.org/abs/2509.09372v2>
- <http://arxiv.org/abs/2110.06500v2>